

PSD-Constrained PAR Reduction for DMT/OFDM

Niklas Andgart

Signal Processing Group, Department of Information Technology, Lund University, P.O. Box 118, 221 00 Lund, Sweden
Email: niklas.andgart@it.lth.se

Brian S. Krongold

Australian Research Council (ARC) Special Research Centre for Ultra-Broadband Information Networks,
University of Melbourne, Victoria 3010, Australia
Email: bsk@ee.mu.oz.au

Per Ödling

Signal Processing Group, Department of Information Technology, Lund University, P.O. Box 118, 221 00 Lund, Sweden
Email: per.odling@it.lth.se

Albin Johansson

Ericsson AB, 126 25 Stockholm, Sweden
Email: albin.johansson@ericsson.com

Per Ola Börjesson

Signal Processing Group, Department of Information Technology, Lund University, P.O. Box 118, 221 00 Lund, Sweden
Email: per.ola.borjesson@it.lth.se

Received 28 February 2003; Revised 30 January 2004

Common to all DMT/OFDM systems is a large peak-to-average ratio (PAR), which can lead to low power efficiency and nonlinear distortion. Tone reservation uses unused or reserved tones to design a peak-canceling signal to lower the PAR of a transmit block. In DMT ADSL systems, the power allocated to these tones may be limited due to crosstalk issues with many users in one twisted pair bundle. This PSD limitation not only limits PAR reduction ability, but also makes the optimization problem more challenging to solve. Extending the recently proposed active set tone reservation method, we develop an efficient algorithm with performance close to the optimal solution.

Keywords and phrases: peak-to-average ratio, tone reservation, OFDM, DMT.

1. INTRODUCTION

Communication systems using multicarrier modulation have recently become widely used both in wireless (DVB-T, DAB, IEEE 802.11a) and wireline (ADSL, VDSL) environments [1, 2, 3]. Multicarrier systems have distinct advantages over single-carrier systems, but suffer from a serious drawback: the approximately Gaussian-distributed output samples cause a high peak-to-average ratio (PAR) that results in low power efficiency and possible nonlinear distortion.

In order to alleviate this PAR problem, many researchers have made efforts to reduce large signal peaks through a variety of PAR reduction methods [4, 5, 6, 7, 8, 9, 10]. A technique known as *tone reservation* was initially developed in [4, 5] and is well suited for discrete multitone modulation (DMT) ADSL systems over twisted pair copper wiring. A

common phenomenon of this environment is a distance-dependent rolloff of the channel transfer function power with increasing frequency, resulting in upper frequency subchannels having very low SNRs and being incapable of reliably transmitting data. An additive peak-canceling signal can be constructed from these dataless tones, as in [4, 5], to help reduce the PAR problem. Further developed tone reservation algorithms have been presented in [11, 12, 13, 14, 15, 16, 17, 18].

In ADSL and other practical systems, the peak-reduction signal may be power limited on each of the reserved tones due to crosstalk constraints with many users being serviced in one twisted pair bundle. This is, for instance, manifested in the recent ADSL2 standard [19] as a -10 dB PSD limit on the reserved tones compared to the data-carrying tones. This PSD constraint on the tones can change the theoretical




FIGURE 1: Addition of a PAR reduction signal, $c[n]$, that counteracts the peaks in $x[n]$. The signal $c[n]$, comprised of a small subset of tones, is a function of the data signal $x[n]$.

ability of tone reservation to reduce the PAR [20, 21] as well as the complexity versus performance tradeoff for practical algorithms.

In this paper, we analyze the PSD-constrained tone reservation problem and its complexity versus performance tradeoff. We extend the recently proposed active set tone reservation approach [16] to handle PSD constraints. Results are analyzed and compared to performance bounds, and computational complexity and algorithm alteration are detailed. In Section 2, we define the system and data model, give a description of the active set PAR reduction algorithm, and introduce PSD-constrained tone reservation. Extension of the active set approach to the PSD-constrained case is presented and analyzed in Section 3, followed by simulation results presented in Section 4.

2. DMT AND TONE RESERVATION

A DMT system uses a symbol length of N samples, which is typically 512 samples in the ADSL downstream direction. Although these samples uniquely define a signal block, when considering the PAR of the analog signal, *peak regrowth* [16, 17, 18] between the sampling points upon digital-to-analog (D/A) conversion has to be considered. Oversampling of the digital signal is a viable approach.

Figure 1 schematically describes the reduction approach. A reduction signal $c[n]$ is added to the original data signal $x[n]$, and is constructed of dataless tones that either cannot transmit data reliably (due to low SNRs) or are explicitly reserved by the system for PAR reduction. For example, in the ADSL2 standard, the mechanism for this is to exclude the reserved tones from the *supported set* of data tones during startup. The goal for the PAR reduction algorithm is to make the resulting signal, $\tilde{x}[n] = x[n] + c[n]$, have a smaller amplitude span than $x[n]$. If the reduction signal is constructed of tones with low SNRs, the reduction signal $c[n]$ may be attenuated before arriving at the receiver. This makes tone reservation using low SNR tones mainly applicable to reducing the transmitter side PAR.

The PAR is defined as

$$\text{PAR}\{\tilde{\mathbf{x}}\} = \frac{\max_n |x[n] + c[n]|^2}{E[|x[n]|^2]}, \quad (1)$$

where the average power in the denominator is that of the

data-bearing signal before PAR reduction is applied.¹ We define

$$\begin{aligned} \tilde{x}[n] &= x[n] + c[n] \\ &= \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} (X_k + C_k) e^{j2\pi kn/N}, \end{aligned} \quad (2)$$

where X_k represents the data symbols and C_k the FFT domain PAR reduction signal. On a given DMT tone, one of them has to be zero to maintain distortionless data transmission

$$X_k + C_k = \begin{cases} X_k, & k \in \mathcal{U}^c, \\ C_k, & k \in \mathcal{U}, \end{cases} \quad (3)$$

where \mathcal{U}^c represents the set of data-bearing subchannels and \mathcal{U} represents the set of available subchannels for PAR reduction.

Let \mathbf{x}_L denote the data signal of one symbol block and let \mathbf{c}_L denote the additive peak-reduction signal generated from the tone set \mathcal{U} , both oversampled to L times the nominal sample rate. We focus on the specific case of a real baseband DMT system, where the data and reduction signals can be expressed as weighted sums of real-valued sinusoids and cosinusoids. In matrix form, we can write $\mathbf{c}_L = \check{\mathbf{Q}}_L \check{\mathbf{C}}$, where $\check{\mathbf{Q}}_L$ is an $NL \times 2U$ matrix of sinusoidal and cosinusoidal column vectors with frequencies specified by the U reserved tones t_1, \dots, t_U ,

$$\check{Q}_{L(i,j)} = \begin{cases} \cos \frac{2\pi(i-1)t_{(j+1)/2}}{NL}, & i \text{ odd}, \\ \sin \frac{2\pi(i-1)t_{j/2}}{NL}, & i \text{ even}, \end{cases} \quad (4)$$

and $\check{\mathbf{C}}$ is a length $2U$ vector with the weights of these (co)sinusoids,

$$\check{C}_{(i)} = \begin{cases} \frac{2}{\sqrt{N}} \text{Re}\{C_{t_{(i+1)/2}}\}, & i \text{ odd}, \\ -\frac{2}{\sqrt{N}} \text{Im}\{C_{t_{i/2}}\}, & i \text{ even}. \end{cases} \quad (5)$$

For this real-valued case, minimizing the peak magnitude of the resulting signal, equivalent to minimizing its peak power, can be formulated as the linear program [5]

$$\begin{aligned} &\text{minimize } \gamma \\ &\text{subject to } \begin{cases} \mathbf{x}_L + \check{\mathbf{Q}}_L \check{\mathbf{C}} \leq \gamma, \\ -\mathbf{x}_L - \check{\mathbf{Q}}_L \check{\mathbf{C}} \leq \gamma. \end{cases} \end{aligned} \quad (6)$$

2.1. Tone selection

It is desired that reduction signal \mathbf{c}_L cancels out the peaks in the data signal \mathbf{x}_L as best as possible. Total cancellation,

¹Although it is mainly referred to as the PAR problem, the real issue is the peak power at the high power amplifier (HPA), in DSL systems commonly called the line driver. Reducing the PAR by inflating the average power does not help. The average power is simply a way of normalizing peak power results, and this normalization factor should remain constant for comparison purposes.

$\mathbf{c}_L = -\mathbf{x}_L$, is naturally impossible, and an alternative, yet still unrealistic, goal is to drive the signal towards a PAR of 0 dB (i.e., the peak power and average power are equal). This tight control of the signal requires a large portion of the frequency band. In general, more reserved tones lead to a lower PAR, and therefore, a tradeoff exists between data throughput and PAR [22]. A choice must be made as to which tones will be used for PAR reduction rather than data transmission. If the system is able to freely choose, the distribution of these tones over the system bandwidth has a significant impact on PAR reduction ability. In general, with no power constraints on the reduction tones, an uneven, spread-out placement (e.g., generated by a random selection of tones) allows for very good PAR reduction [5, 23]. A significant performance loss, however, results by placing the reduction tones as a contiguous block or uniformly distributed over the entire bandwidth.

In wireline DMT systems, it is preferred to use those tones which cannot send data reliably due to insufficient SNRs, thereby maintaining the same throughput level. Generally, these tones are in the uppermost frequencies, and tend to resemble a contiguous block of tones, which is not a good tone set in terms of performance. An alternative is to reduce the system throughput by sacrificing some tones for peak reduction and achieving an uneven, spread-out placement. We will consider these two extreme cases of tone placement. In practice, a combination of these may turn out to be the most attractive choice.

After determining the set of reserved tones, the reduction signal \mathbf{c}_L is created from a nominal peak-reduction kernel \mathbf{p} [5], formed by projecting an impulse at $n = 0$ onto the set of tones \mathcal{U} . This corresponds to the least squares approximation of the impulse with equal weight on each reduction tone. Other forms of \mathbf{p} generated by different criteria, such as minimizing the size of their sidelobes, have been suggested in [5].

2.2. Active set tone reservation

The linear program in (6) can be solved with a simplex method, but is expensive with a complexity of $O(N^2L^2)$ operations. Computationally efficient $O(NL)$ approaches based upon projection-onto convex sets (POCS) and gradient projection were developed in [4, 5], respectively, but suffer from slow convergence. A recent $O(NL)$ approach [13, 16] was developed based on active set methods [24] and exhibits rapid convergence towards a minimax PAR solution. Whereas a finite number of iterations will achieve the optimal PAR level γ^* for the given tone set, a very good suboptimal solution can be achieved in two or three iterations, making this an attractive practical solution.

As in the gradient project and POCS approaches, the active set approach reduces the PAR through the use of the kernel \mathbf{p} . Circularly shifted versions of this kernel, $\mathbf{p}_{(\cdot)}$, also lie in the signal space generated from \mathcal{U} , allowing easy reduction of a peak at an arbitrary sample location.

Beginning with the sample of largest magnitude γ_0 at location n_0 , the peak is reduced by subtracting a scaled version of $\mathbf{p}_{(n_0)}$ until a second peak at some location n_1 is *balanced*

with it at some magnitude $\gamma_1 < \gamma_0$. These two peaks are then reduced equally through a linear combination of $\mathbf{p}_{(n_0)}$ and $\mathbf{p}_{(n_1)}$ until a third peak is balanced. These three peaks are reduced equally until a fourth is balanced, and so forth. When a sample is at the peak magnitude, it signifies an active inequality constraint (i.e., strictly equal) in (6), and the active set approach is therefore building a set of active constraints. Mathematically, the iteration updates can be written as

$$\bar{\mathbf{x}}^{(i)} = \bar{\mathbf{x}}^{(i-1)} - \mu^{(i)} \hat{\mathbf{p}}^{(i)}, \quad (7)$$

where $\bar{\mathbf{x}}^{(i)}$ represents the signal after the i th iteration, $\hat{\mathbf{p}}^{(i)}$ is the descent direction in the i th iteration, and $\mu^{(i)}$ represents the descent step size.

At the start of the i th iteration, there will be i peaks which are balanced at locations n_0, n_1, \dots, n_{i-1} . To keep these peaks balanced, the next iteration descent must satisfy

$$\hat{p}_{n_k}^{(i)} = \text{sign}(\bar{x}_{n_k}^{(i)}) = S_{n_k}, \quad k = 0, 1, \dots, i-1, \quad (8)$$

with the assumption that we scale $\hat{\mathbf{p}}^{(i)}$ to have unit magnitude in locations corresponding to the active set of peaks. No matter what value of $\mu^{(i)}$ is chosen, the magnitudes of the peaks at n_0, n_1, \dots, n_{i-1} will remain equal. The \hat{p}_{n_i} values can be calculated as

$$\hat{\mathbf{p}}^{(i)} = \sum_{k=0}^{i-1} \alpha_k^{(i)} \mathbf{p}_{(n_k)}, \quad (9)$$

where the $\alpha_k^{(i)}$ are computed by solving the $i \times i$ system of equations

$$\begin{bmatrix} 1 & p_{n_0-n_1} & \cdots & p_{n_0-n_{i-1}} \\ p_{n_1-n_0} & 1 & \cdots & p_{n_1-n_{i-1}} \\ \vdots & \vdots & \ddots & \vdots \\ p_{n_{i-1}-n_0} & p_{n_{i-1}-n_1} & \cdots & 1 \end{bmatrix} \begin{bmatrix} \alpha_0^{(i)} \\ \alpha_1^{(i)} \\ \vdots \\ \alpha_{i-1}^{(i)} \end{bmatrix} = \begin{bmatrix} S_{n_0} \\ S_{n_1} \\ \vdots \\ S_{n_{i-1}} \end{bmatrix}. \quad (10)$$

This requires an $i \times i$ matrix inverse, but in practical implementations, i will typically be at most 3, and the inverse cost is then insignificant relative to the total iteration complexity. Furthermore, efficient inverse techniques [25] can be applied as in addition to being symmetric (due to the symmetry of \mathbf{p}), the matrix in a given iteration is contained in the matrix for the next iteration.

The step size $\mu^{(i)}$ required to balance the next active peak is determined by testing samples as follows² (see [15, 16] for more details),

$$\mu^{(i)} = \min_{q \in \mathcal{A}} \left(\frac{\gamma^{(i-1)} - |\bar{x}_q^{(i-1)}|}{1 - \text{sign}(\bar{x}_q^{(i-1)}) \hat{p}_q^{(i)}} \geq 0 \right), \quad (11)$$

where \mathcal{A} represents the set of samples in the active set. Strategies exist [15, 16] to reduce the sample testing complexity

²The $\min(\cdot \geq 0)$ notation means to take the minimum over the nonnegative elements.

as the structure of $\bar{\mathbf{x}}^{(i)}$ and $\hat{\mathbf{p}}^{(i)}$ can be exploited to eliminate many potential samples from consideration. For practical implementation, the division operation can be replaced by a multiplication with the output of a prestored inverse lookup table to approximate $1/(1 - \text{sign}(\bar{x}_q^{(i-1)})\hat{p}_q^{(i)})$. Exact values are not needed for comparison purposes, and therefore, a dense lookup table is not required.

2.3. PSD-constrained tone reservation

Solving (6) for the optimum PAR value will in many cases cause the power on the reduction tones to grow immensely as the very last bits of reduction performance require large reduction signals. A standardized system generally has to follow certain PSD constraints on data tones. Similar rules are applicable for reduction tones as well, especially in wireline systems where crosstalk exists and the effect on other users should be kept to a minimum. Thus, a system may have to abide by instantaneous and/or average power constraints on the reserved tones.

What the PSD constraint should be is a system design issue based upon factors such as crosstalk and power consumption or, in practice, often determined by a standard. In the new ADSL2 ITU-T Recommendation [19, Figure 8-19/G.992.3], *passband* tones are under strict control and can be grouped into different categories: one group of tones is for data transmission and another group consists of monitored tones for receiver functions (e.g., channel estimation). Both of these groups belong to the *medley set*. Tones that are not in the medley set have a PSD restriction 10 dB below the nominal PSD level and these are the tones that can be used for PAR reduction.

Since the PSD is a measurement averaged over time, the power on the tones may be allowed to vary from symbol to symbol, and the instantaneous power of a symbol may therefore exceed the PSD constraint. As an example, consider a target PAR value of 12 dB and the uppermost probability curves for unreduced signals shown in Figures 4 to 9. It follows that approximately 8% of the symbols require PAR reduction, and due to averaging, a revised PSD constraint on the reserved tones can be determined. If PAR reduction is employed for only 8% of the symbols, we can allow an average reserved tone power $10 \log(1/0.08) \approx 11$ dB above its overall -10 dB PSD constraint. This results in a revised PSD constraint on the reserved tones $-10 \text{ dB} + 11 \text{ dB} = +1 \text{ dB}$ above the nominal PSD mask for the ADSL2 system.

When processing one symbol at a time, however, a peak power constraint per tone for each symbol is much easier to deal with than an averaged PSD constraint. Using this power constraint can cause the averaged PSD figure to be somewhat less than this peak constraint. Nevertheless, for a given peak power constraint per tone, a corresponding averaged PSD level can be determined experimentally for a specific system, and the constraints can then be interchanged. In the rest of this paper, we consider the peak power limitation, or *instantaneous* PSD constraint, on each tone rather than a PSD as a result of averaging.

Incorporating the power constraint on each tone, the PSD restriction becomes part of (6) in the form of a quad-

ratic constraint:

$$\begin{aligned} & \text{minimize } \gamma \\ & \text{subject to } \begin{cases} \mathbf{x}_L + \check{\mathbf{Q}}_L \check{\mathbf{C}} \leq \gamma, \\ -\mathbf{x}_L - \check{\mathbf{Q}}_L \check{\mathbf{C}} \leq \gamma, \\ \check{\mathbf{C}}_{(2l-1)}^2 + \check{\mathbf{C}}_{(2l)}^2 \leq A_{l,\max}^2, \end{cases} \end{aligned} \quad (12)$$

where $A_{l,\max}$ is the limitation in amplitude on tone l . Due to the introduction of quadratic constraints, the problem is no longer a linear program, but instead a *quadratically constrained quadratic program* (QCQP).

3. PSD-CONSTRAINED ACTIVE SET APPROACH

3.1. Modifications for PSD constraints

If the active set algorithm is to be used in the PSD-constrained case, it must be modified. Letting \check{C}_l denote the l th element of $\check{\mathbf{C}}$ (including both cosine and sine parts), the total weight on tone l after iteration i can be described as

$$\check{C}_l^{(i)} = \check{C}_l^{(i-1)} + \Delta \check{C}_l^{(i)}, \quad (13)$$

where the increments $\Delta \check{C}_l^{(i)}$ in each iteration include the effect from reducing one additional peak. Using the step size $\mu^{(i)}$ and weighting $\alpha_k^{(i)}$ from (9), the increments $\Delta \check{C}_l^{(i)}$ can be expressed in cosine and sine components.

$$\begin{aligned} \Delta \check{C}_l^{(i)} &= \begin{bmatrix} \Delta \check{C}_{l,\cos}^{(i)} & \Delta \check{C}_{l,\sin}^{(i)} \end{bmatrix} \\ &= K \mu^{(i)} \sum_{k=0}^{i-1} \alpha_k^{(i)} \begin{bmatrix} \cos\left(\frac{2\pi t_l n_k}{NL}\right) & \sin\left(\frac{2\pi t_l n_k}{NL}\right) \end{bmatrix}, \end{aligned} \quad (14)$$

where K is a known constant that results from normalizing \mathbf{p} so that $p_0 = 1$. We can think of three main outcomes when performing an active set iteration at an instance where none of the PSD constraints have been met or exceeded.

- (1) *A new peak is balanced and no PSD constraints are met.* This is the same case as with no PSD constraint. The algorithm can continue with its next step.
- (2) *All tones meet/exceed the PSD constraints at the same time.* This happens when reducing one peak and reaching the PSD constraint before a second active peak is encountered.
- (3) *Some tones meet/exceed the PSD constraint.* This can happen when two or more peaks are already balanced. Then different tones will likely have different magnitudes, see Figure 2.

For case (1), the algorithm will be identical to what is described in Section 2.2. For case (2), the algorithm merely takes the step μ_{\max} that fills all subchannels to the PSD constraint, and the optimal solution has been reached.

The interesting question is what to do in case (3), as some of the tones have filled up or gone past their PSD constraints, while others are still available for further reduction. The μ descent can easily be scaled back to where the first tone reaches




FIGURE 2: Addition of the tone weights for reduction of two different peaks can cause the PSD constraint to be reached on certain tones before others.

the PSD constraint, that tone can be frozen, and the remaining tones can be used for PAR reduction for subsequent iterations. This process can be repeated until all tones reach the PSD constraint. We note that an iteration now refers to the operations performed to reach either a new active peak or a new tone that meets the PSD constraint.

3.2. Cost-versus-performance issues

It can be expected that once any tone reaches the PSD constraint, many or all of the remaining tones are not far from reaching it as well. At this point, the problem is that convergence speed (i.e., additional PAR reduction per iteration) is severely reduced as a new iteration must be performed to the point where either a new tone reaches the PSD constraint or a new active peak is encountered.

After each new tone reaches the PSD constraint and is shut off, the set \mathcal{U} changes and a new nominal peak-reduction kernel \mathbf{p} needs to be recomputed. Rather than compute the projection of an impulse onto the remaining tones, the contribution of the removed tone can just be subtracted (using NL operations) from the latest \mathbf{p} .

3.3. Low complexity algorithm

The cost-versus-performance tradeoff dictates that it may not be worth iterating beyond the point where the first tone reaches the PSD constraint, and therefore not utilizing the available remaining power in the other tones. This low complexity approach saves a lot of computation and results in only a small performance loss from the optimal solution as simulations show in the next section. The complexity of this extended algorithm is the same as the unconstrained active set approach with an additional extra cost of keeping track of the signal power in each tone. This cost is insignificant compared to the rest of the algorithm since $U \ll NL$.

During each iteration, a new $\hat{\mathbf{p}}^{(i)}$ is created according to (9), and in parallel to that, the new signal in each tone is calculated, based on the additional contributions according to (14). Before applying (7) and potentially wasting operations, $2U$ multiplies and U adds are used to check the tones powers against the PSD constraint. If any of the tones exceeds the PSD constraint, $\mu^{(i)}$ must be scaled back to find the point where the PSD constraint is met with one or more tones. The




FIGURE 3: Linear approximation of the quadratic magnitude constraints. An octagon is shown here, but a polygon with a larger number of sides can be used for a better approximation.

quadratic equation

$$\left| C_l^{(i-1)} + \beta_l \mu^{(i)} \Delta C_l^{(i)} \right|^2 = A_{l,\max}^2 \quad (15)$$

is solved for β_l for the tone(s) exceeding the PSD constraint, and the minimum β_l value is chosen to scale $\mu^{(i)}$. This modified step size is then used in (7) to compute the final PAR-reduced signal.

3.4. Performance bounds

It is important to gauge how much performance is lost when using this low complexity algorithm that halts PAR reduction once any tone reaches the PSD constraint. Three lower bounds on achievable PAR level are now presented.

3.4.1. Bound on minimum PAR

The resulting PAR level after the low complexity algorithm can be compared to the optimal solution of (12). This equation represents a QCQP, and still is a convex problem. Linear approximations of the quadratic constraint (see Figure 3) can be employed to transform the problem back to linear programming form [21], in order to solve the problem with linear programming algorithms. Thereby, a performance bound³ can be computed through simulations. It should be noted that this bound on the optimal solution is extremely tight when used with polygons of 16 sides and larger.

3.4.2. A_{\max} bound

The constraint on maximum power per tone (equivalent to a constraint on the maximum magnitude) results in limiting the magnitude of the peak-reduction signal to A_{\max} , where

$$A_{\max} = \sum_{l=1}^U A_{l,\max}. \quad (16)$$

We assume that an arbitrary peak-reduction signal can be created, with the only limitation being that its amplitude is

³The polygonal approximation completely bounds the circle, thus, the resulting performance will be at least as good as the quadratically constrained optimal solution.

between $-A_{\max}$ and A_{\max} . As a result, starting with a symbol with peak level $\max |x_L[n]|$, the peak level can at best be reduced down to $\max |x_L[n]| - A_{\max}$. Since this model admits additional degrees of freedom compared to the true reduction signal, it serves as a lower bound on the achievable PAR level. Given a peak value for a symbol block, this can be expressed as

$$\max |\bar{x}_L| = \max |x_L + c_L| \geq \max |x_L| - A_{\max}. \quad (17)$$

This A_{\max} bound shows that when the PSD constraint is quite restrictive and only a small number of tones are reserved, PAR reduction performance is severely limited, even with an arbitrary choice of reduction tones [20, 21]. In this case, a choice of tones discarding the minimum amount of data capacity may be the most favorable.

3.4.3. 2-Bound

The A_{\max} bound from (17) corresponds to the achieved peak level when all tones are filled in order to reduce the largest peak in x_L . A similar bound can be computed after the active set approach has already performed its first iteration. The two balanced peaks can be reduced (without any regard for the other samples, and thus making a bound) until all tones meet the PSD constraint. This bound, which we refer to as the *2-bound*, is simple to simulate because α_0 and α_1 must be of equal magnitude due to the symmetry of \mathbf{p} .

4. SIMULATIONS

A DMT system with symbol length $N = 512$ is simulated with tones 33–255 used for either data transmission or PAR reduction (these system parameters are the same for downlink ADSL transmission). Each of the data-carrying tones uses a 1024-point QAM constellation. Before active set processing, the signals have been oversampled by the factor $L = 4$ to limit analog peak-regrowth effects upon digital-to-analog conversion. It has been observed that operating on the digital $L = 1$ signal does not provide any worthwhile PAR reduction performance at the analog signal [15]. Oversampling to $L = 4$ makes the computational cost increase by a factor of 4, although $L = 2$ could be employed for a performance decrease which varies based upon the number of tones, their locations, and PSD constraints.

As described in Section 2.3, the averaged PSD constraint for the reduction tones could be set to about 1 dB above the nominal PSD mask for the given example. We now use this figure as a guideline for the instantaneous PSD mask in the following simulations. To illustrate the effects when varying the maximum reduction power per tone, the simulations will first use a restrictive constraint set at the nominal PSD mask, and then use a looser mask, where the magnitude is increased by 50% (+3.5 dB).

We view the forthcoming PAR results on a per-symbol basis using the simulated probability that at least one sample in a symbol block exceeds a certain PAR level. This corre-




FIGURE 4: Symbol clip probability for 12 PAR reduction tones, chosen as a contiguous block of the highest tones. Up to four active set iterations are applied, but the algorithm stops once any tone hits the PSD constraint. The three leftmost curves represent optimal solution bounds.

sponds to taking the maximum value over one symbol in (1), thereby reflecting the probability that a symbol is transmitted with distortion. This clip probability also is commonly used in the literature. A viable alternative would be to evaluate the clip probability of each individual sample, which reflects the percentage of time the transmitted signal is clipped.

4.1. Block placed tones

4.1.1. Restrictive PSD constraint

Figure 4 shows simulations with the upper block of 12 tones (number 244–255) used for PAR reduction and subjected to an instantaneous PSD constraint equal to the nominal PSD level for the data tones. The curves show the reduction performance using the extended active set algorithm, stopping as soon as any PSD constraint is reached. Shown on the vertical axis is the probability that the time domain symbol block \bar{x}_L would be clipped if subjected to a clip level γ_c on the x -axis, that is,

$$\text{Prob}(\text{PAR}\{\bar{x}_L\} > \gamma_c). \quad (18)$$

Starting at the rightmost line, corresponding to the clip probability of an unreduced symbol, curves representing iterations one through four are shown.

The two leftmost curves show the lower bounds from Section 3.4 (A_{\max} bound and 2-bound), which the simulations cannot cross. The third lowest curve, dashed and

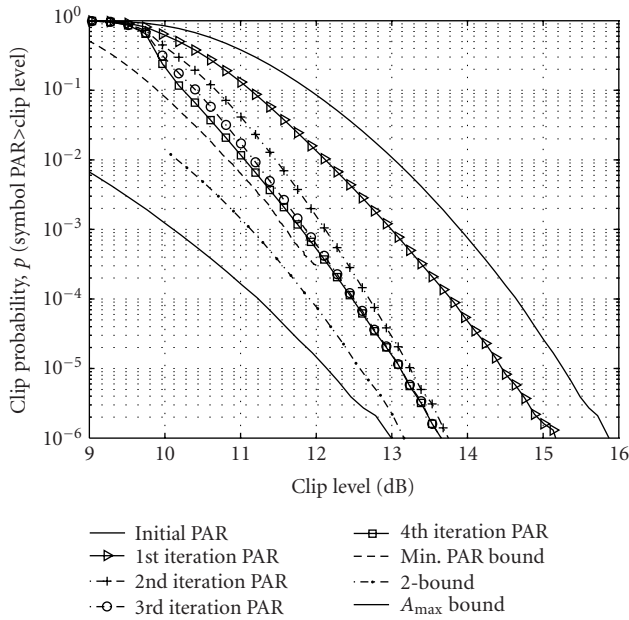


FIGURE 5: Symbol clip probability for PAR reduction with the 12 highest tones. The PSD constraint allows 50% higher magnitude per tone than in Figure 4. The reduction performance shows only a small gain compared to Figure 4, showing that this placement cannot take much advantage of the loosened PSD constraint.

ending at a clip probability of $3 \cdot 10^{-4}$ is the PAR achieved by finding the minimum value of (6) with linearized quadratic constraints (a 32-sided polygon, cf. Figure 3) and using the same upper block of 12 tones. This curve will also serve as a bound for the suboptimal algorithm, but due to its much larger complexity, this curve has not been simulated for the lower clip probabilities.

Looking at the performance of the low complexity algorithm, we see that for the higher clip probabilities, there is a performance gain of about 0.15 dB going beyond two iterations, and an additional 0.1 dB compared to the minimum PAR bound (dashed line). At the lower clip probabilities, we see that the curves converge towards the A_{\max} bound from (17).

Here we see a situation where a restrictive PSD constraint and a small number of reduction tones set a limit on the achievable PAR level. The reduction performance is limited by the A_{\max} bound, and not necessarily by the block placement reduction performance. The low complexity algorithm provides near-optimal performance at a very low cost for this system.

4.1.2. Loosening the PSD constraint

In Figure 5, the PSD constraint is increased by 50% in magnitude for each tone. Comparing the figures, we see that the lower bound decreases due to an increase of the maximum reduction signal. However, the simulated reduction performance, including the optimal solution, increases by




FIGURE 6: Symbol clip probability for PAR reduction with the 24 highest tones with the same PSD constraint used in Figure 4. The simulations indicate only a small reduction gain compared to Figure 4, showing that adding extra tones to the reserved block does not help PAR reduction much.

only about 0.3 dB. The block placement simply cannot take advantage of the increased reduction power, and is the real limiting factor in this case. Looking at the performance of the low complexity algorithm, we see that its loss compared to the minimum PAR bound is about 0.2 dB.

4.1.3. Increasing the number of tones

Figure 6 shows results for when the upper block of 24 tones are used for PAR reduction along with the same PSD constraint as in Figure 4. Looking at the figure, we see that the gain from 12 to 24 tones is only about 0.4 dB, which is small considering that the maximum reduction magnitude has been doubled (the A_{\max} bound is significantly lower). In this situation, however, we see that after 4 active set iterations, we are about 0.2 dB from the minimum PAR bound at higher probabilities, thus telling us that further iterations are likely not worth the significant cost to achieve it.

4.2. Randomly chosen tones

We have seen that even when constraints (PSD limit or number of tones) are loosened, a bad tone set selection can still be a limiting factor. Now a more “spread-out” toneset is evaluated, where the reserved tones are randomly selected in the interval from 33 to 255 inclusive.

4.2.1. Restrictive PSD constraint

Figure 7 shows similar simulations as Figure 4, using the restrictive instantaneous PSD constraint, equal to the average




FIGURE 7: Symbol clip probability for 12 randomly chosen PAR reduction tones. The three lowest curves show bounds on the achievable performance as in previous simulations.




FIGURE 8: Symbol clip probability for PAR reduction with 12 random tones. The PSD constraint allows 50% higher magnitude per tone than in Figure 7.

power mask for the data tones. Looking at the figure, the iterations quickly converge to within 0.1 dB of the A_{max} bound, and the performance is only slightly better than for block placed tones. Here the A_{max} bound effectively sets the limitation on system performance [20, 21].




FIGURE 9: Symbol clip probability for PAR reduction with 24 random tones with the same PSD constraint used in Figure 7.

4.2.2. Loosening the PSD constraint

Figure 8 shows the performance when the PSD constraint is set to allow for a tone magnitude 50% higher than before. The reduction performance has increased thanks to more allowed power. At the lower clip probabilities, the gains are close to 1 dB compared to Figure 7, and the active set results are very close to the performance bounds. At higher clip probabilities, the gains are close to 0.5 dB, but are a noticeable distance from the very tight minimum-PAR bound. This is only a minor issue, since in these regions, the PAR level after 3 or 4 iterations is already rather low.

4.2.3. Increasing the number of tones

Finally, Figure 9 shows simulations using 24 randomly chosen tones, with the restrictive PSD constraint. Due to the superior reduction ability for this placement type, the resulting PAR level is clearly lower than in the previous simulation. The allowed A_{max} is 100% higher here than with half the number of tones, and we see that a larger number of active set iterations may be needed to achieve PAR levels very close to the optimal solution. However, when considering lower clip probabilities, the 4th active set iteration is not very far from the 2-bound.

5. CONCLUSIONS

Introducing PSD constraints into tone reservation affects the achievable PAR reduction and significantly alters the complexity-versus-performance tradeoff for practical algorithms. The results in this paper show the impact that PSD constraints have on tone reservation performance, and it is clear

that the effect when using randomly chosen tone sets is more severe than for contiguous tone sets.

A low complexity suboptimal solution has been presented, and results show that its performance is close to optimal solution bounds. Since small performance increases incur a major computation cost (greater than the low complexity algorithm itself), we assert that our proposed approach gives a very good tradeoff of complexity and PAR reduction.

To evaluate whether the oversampling of $L = 4$ is sufficient, the signals were oversampled by an additional factor of 4 after reduction. The peak regrowth has been observed to be less than 0.2 dB. Further studies could also include the effect on peak regrowth after the filter chain present in the transmitter [16, 17, 18].

An important special case results when a nonuniform PSD constraint is given, that is, more power is allowed on some reserved tones than others. In this case, certain tones may reach their PSD constraint much sooner than the rest, and sizeable performance gains beyond this stoppage point may still exist. An intelligent approach may be to modify the formation of \mathbf{p} by weighting the impulse projection onto the tones according to the nonuniformity of the PSD mask. In this way, the more restricted tones do not reach their PSD constraint with greater ease than the others.

Although the real baseband DMT case is the main focus of this paper, the principles can also be applied to the complex baseband case (for wireless OFDM systems), as an active set approach for this case has already been developed in [14, 16]. The problem with tone reservation in wireless systems is that it may not be desirable to sacrifice data tones in a fading channel. However, it is possible that in a fixed wireless scenario (with a slowly varying channel), channel state feedback could be employed and certain subchannels with low SNRs could be used for tone reservation.

ACKNOWLEDGMENT

This work was supported by Ericsson AB and by the Australian Research Council.

REFERENCES

- [1] J. A. C. Bingham, "Multicarrier modulation for data transmission: an idea whose time has come," *IEEE Communications Magazine*, vol. 28, no. 5, pp. 5–14, 1990.
- [2] T. Starr, J. M. Cioffi, and P. J. Silverman, *Understanding Digital Subscriber Line Technology*, Prentice Hall, Upper Saddle River, NJ, USA, 1999.
- [3] ITU-T, *Asymmetric digital subscriber line (ADSL) transceivers*, Recommendation G.992.1, June 1999.
- [4] A. Gatherer and M. Polley, "Controlling clipping probability in DMT transmission," in *Proc. Asilomar Conference on Signals, Systems, and Computers*, vol. 1, pp. 578–584, Pacific Grove, Calif, USA, November 1997.
- [5] J. Tellado-Mourelo, *Peak to average power reduction for multicarrier modulation*, Ph.D. dissertation, Stanford University, Stanford, Calif, USA, September 1999.
- [6] M. Friese, "Multitone signals with low crest factor," *IEEE Trans. Communications*, vol. 45, no. 10, pp. 1338–1344, 1997.
- [7] D. L. Jones, "Peak power reduction in OFDM and DMT via active channel modification," in *Proc. Asilomar Conference on Signals, Systems, and Computers*, vol. 2, pp. 1076–1079, Pacific Grove, Calif, USA, October 1999.
- [8] D. J. G. Mestdagh and P. M. P. Spruyt, "A method to reduce the probability of clipping in DMT-based transceivers," *IEEE Trans. Communications*, vol. 44, no. 10, pp. 1234–1238, 1996.
- [9] B. M. Popović, "Synthesis of power efficient multitone signals with flat amplitude spectrum," *IEEE Trans. Communications*, vol. 39, no. 7, pp. 1031–1033, 1991.
- [10] R. W. Bäuml, R. F. H. Fischer, and J. B. Huber, "Reducing the peak-to-average power ratio of multicarrier modulation by selected mapping," *Electronics Letters*, vol. 32, no. 22, pp. 2056–2057, 1996.
- [11] P. O. Börjesson, H. G. Feichtinger, N. Grip, et al., "A low-complexity PAR-reduction method for DMT-VDSL," in *Proc. 5th IEEE International Symposium on Digital Signal Processing for Communication Systems*, pp. 164–199, Perth, Australia, February 1999.
- [12] P. O. Börjesson, H. G. Feichtinger, N. Grip, et al., "DMT PAR-reduction by weighted cancellation waveforms," in *Proc. Radioteknisk Konferens*, pp. 303–307, Karlskrona, Sweden, June 1999.
- [13] B. S. Krongold and D. L. Jones, "A new method for PAR reduction in baseband DMT systems," in *Proc. Asilomar Conference on Signals, Systems, and Computers*, vol. 1, pp. 502–506, Pacific Grove, Calif, USA, November 2001.
- [14] B. S. Krongold and D. L. Jones, "A new tone reservation method for complex-baseband PAR reduction in OFDM systems," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 3, pp. 2321–2324, Orlando, Fla, USA, May 2002.
- [15] B. S. Krongold, *New techniques for multicarrier communication systems*, Ph.D. dissertation, University of Illinois at Urbana-Champaign, Urbana, Ill, USA, November 2001.
- [16] B. S. Krongold and D. L. Jones, "An active-set approach for OFDM PAR reduction via tone reservation," *IEEE Trans. Signal Processing*, vol. 52, no. 2, pp. 495–509, 2004.
- [17] W. Henkel and V. Zrno, "PAR reduction revisited: an extension to Tellado's method," in *Proc. International OFDM-Workshop*, pp. 31.1–31.6, Hamburg, Germany, September 2001.
- [18] J. Tellado and J. M. Cioffi, *Further results on peak-to-average ratio reduction*, ANSI Document, T1E1.4 no. 98-252, August 1998.
- [19] ITU-T, *Asymmetric digital subscriber line (ADSL) transceivers-2 (ADSL2)*, Recommendation G.992.3, July 2002.
- [20] N. Petersson, A. Johansson, P. Ödling, and P. O. Börjesson, "A performance bound on PSD-constrained PAR reduction," in *Proc. IEEE International Conference on Communications*, pp. 3498–3502, Anchorage, Alaska, USA, May 2003.
- [21] N. Petersson, *Peak and power reduction in multicarrier systems*, Licentiate thesis, Lund University, Lund, Sweden, November 2002.
- [22] P. Ödling, N. Petersson, A. Johansson, and P. O. Börjesson, "How much PAR to bring to the party?," in *Proc. Nordic Signal Processing Symposium*, Tromsø-Trondheim, Norway, October 2002.
- [23] N. Petersson, A. Johansson, P. Ödling, and P. O. Börjesson, "Analysis of tone selection for PAR reduction," in *Proc. International Conference on Information, Communications and Signal Processing*, Singapore, October 2001.
- [24] D. G. Luenberger, *Linear and Nonlinear Programming*, Addison-Wesley, Boston, Mass, USA, 1984.
- [25] G. H. Golub and C. F. van Loan, *Matrix Computations*, John Hopkins University Press, Baltimore, Md, USA, 2nd edition, 1989.

Niklas Andgart was born in Hässleholm, Sweden in 1975. He received his M.S.E.E. degree in 2000 and his Licentiate in Engineering degree in 2002, both from Lund University. During the fall of 1999 he was with the Vehicle and Dynamics Laboratory at the University of California at Berkeley, and in early 2004 he was visiting the Department of Electrical and Electronic Engineering at the University of Melbourne. Currently, he is working towards a Ph.D. in signal processing at the Department of Information Technology at Lund University. His research is within signal processing for communication systems and he works with DSL research in cooperation with Ericsson AB in Stockholm.



Brian S. Krongold received his B.S., M.S., and Ph.D. degrees in electrical engineering in 1995, 1997, and 2001, respectively, from the University of Illinois at Urbana-Champaign, and worked as a Research Assistant at the Coordinated Science Laboratory from 1995–2001. Since December 2001, he has been a Research Fellow in the ARC Special Research Centre for Ultra-Broadband Information Networks in the Department of Electrical and Electronic Engineering at the University of Melbourne, Australia. During the summer of 1994, he interned for Martin Marietta at the Oak Ridge National Laboratory, Oak Ridge, Tennessee. From January to August 1995, he consulted at Bell Laboratories in Middletown, New Jersey. During the summer of 1998, he worked at the Electronics and Telecommunications Research Institute, Taejon, South Korea, under a National Science Foundation summer research grant. He received his second prize in the Student Paper Contest at the 2001 Asilomar Conference on Signals, Systems, and Computers. His research interests are in multicarrier communication systems, electro-optical signal processing, and time-frequency analysis and wavelets.



Per Ödling was born in 1966 in Örnsköldsvik, Sweden. He received his M.S.E.E. degree in 1989, his Licentiate of Engineering 1993, and his Ph.D. in signal processing 1995, all from Luleå University of Technology, Sweden. In 2000, he was awarded the Docent degree from Lund Institute of Technology, and in 2003 he was appointed Full Professor there. From 1995, he was an Assistant Professor at Luleå University of Technology, serving as Vice Head of the Division of Signal Processing. In parallel, he consulted for Telia AB and ST-Microelectronics, developing an OFDM-based proposal for the standardization of UMTS/IMT-2000 and VDSL for standardization in ITU, ETSI, and ANSI. Accepting a position as Key Researcher at the Telecommunications Research Center, Vienna in 1999, he left the arctic north for historic Vienna. There, he spent three years advising graduate students and industry. He also consulted for the Austrian Telecommunications Regulatory Authority on the unbundling of the local loop. He is, since 2003, a Professor at Lund Institute of Technology, stationed at Ericsson AB, Stockholm. He also serves as an Associate Editor for the IEEE Transactions on Vehicular Technology. He has published more than forty journal and conference papers, thirty-five standardization contributions, and a dozen patents.

Albin Johansson was born in 1968 in Stockholm, Sweden. He received his M.S.E.E. degree in 1993 from Royal Institute of Technology in Stockholm and is now pursuing his Ph.D. at Lund Institute of Technology. From 1993 he holds a position at Ericsson AB as Chief of Technology Linecards within broadband access, being responsible for the choice of the technology in Ericsson's wireline broadband access products. He has been actively involved in development of the standardization of ADSL within ETSI, ANSI, ITU-T, and ADSL forum. He has been Editor for ITU-T G.997.1 and chair in one of ADSL forums subcommittees. In addition, from 1992 to 1995, he was teaching undergraduate students at Royal Institute of Technology. Since 2001, he has been a member of the Signal Processing group at Lund Institute of Technology. He has published 6 conference papers, numerous standardization contributions, and holds 7 patents.

Per Ola Börjesson was born in Karlshamn, Sweden in 1945. He received his M.S. degree in electrical engineering in 1970 and his Ph.D. degree in telecommunication theory in 1980, both from Lund Institute of Technology (LTH), Lund, Sweden. In 1983, he received the degree of Docent in telecommunication theory. From 1988 to 1998, he was Professor of signal processing at Luleå University of Technology. Since 1998, he is Professor of signal processing at Lund University. His primary research interest lies in high performance communication systems, in particular, high data rate wireless and twisted pair systems. He is presently researching signal processing techniques in communication systems that use orthogonal frequency division multiplexing (OFDM) or discrete multitone (DMT) modulation. He emphasizes the interaction between models and real systems, from the creation of application-oriented models based on system knowledge to the implementation and evaluation of algorithms.

